NEIL PROTOCOL ENGINE

HYPER-DETAILED BOOK DECONSTRUCTION

THE ARTILECT WAR: COSMISTS VS. TERRANS

Hugo de Garis

PART I: FOUNDATIONAL OVERVIEW

Bibliographic Profile

Full title: The Artilect War: Cosmists vs. Terrans — A Bitter Controversy Concerning Whether Humanity Should Build Godlike Massively Intelligent Machines

Author: Hugo de Garis, Ph.D.

Publication year: 2005 (ETC Publications)

Subject area: Futurism, philosophy of artificial intelligence, speculative geopolitics

Intellectual tradition: Transhumanist-adjacent; technological singularity theory; evolutionary computation; speculative political philosophy

Historical Context of Publication

The book appeared in the early years of what became known as the AI spring's warm-up period — post-dot-com crash, pre-deep-learning revolution. The broader cultural moment included Ray Kurzweil's The Singularity Is Near (2005), Nick Bostrom's early existential risk papers, and sustained media attention to AI futures following the plateau of 1990s AI hype. De Garis was writing in a window when serious institutional AI alignment work barely existed as a field. The Unabomber's manifesto was a decade old. The War on Terror had just begun. Moore's Law still functioned as a near-religious

doctrine in Silicon Valley. Against this backdrop, The Artilect War reads as a book that arrived slightly before its natural audience existed — a condition de Garis explicitly acknowledged, framing himself as a Szilard-type figure warning about a threat the establishment was not yet equipped to process.

The book was ignored by mainstream academic philosophy and received limited attention from AI researchers. The communities that engaged with it were early transhumanist networks, futurist conferences, and the embryonic existential risk community that would later coalesce around organizations like MIRI and FHI. Its ideas were effectively absorbed into that discourse without the book itself becoming canonical. De Garis remained a peripheral figure — his technical credentials real but his institutional standing modest, his prose too raw for philosophy journals and too philosophical for engineering ones.

The Author's Background and Its Effect on the Text

De Garis was born in Sydney in 1947, educated in applied mathematics and theoretical physics, and spent his career moving between countries — UK, Belgium, Japan, USA, China — which gave him a genuine outsider's perspective on any single culture's assumptions about technology and progress. He pioneered the field of evolvable hardware and built what he called CAM-Brain Machines (CBMs), hardware systems capable of evolving neural net modules using genetic algorithms at electronic speeds. He claimed co-authorship of the research fields of evolutionary neural systems and evolvable hardware. The Guinness Book of World Records listed his CBM. He is a credible technologist of the second tier — not a household name in AI, but not a crank. His celebrity was largest in France and the Netherlands, where he received senate-level political attention.

The autobiographical chapter of the book is unusually candid. De Garis describes multiple marriages, a wife killed by lung cancer (her mother was gassed at Auschwitz), financial loss in the dot-com crash, personal schizophrenia about his work, recurring nightmares about his grandchildren dying in a gigadeath war. This confessional texture is not incidental. The book is partly a public confession from someone who believes he is helping to construct the conditions for civilizational catastrophe and cannot stop himself. The Faustian self-awareness is the book's most honest register.

Central Thesis

The book's argument can be compressed into three claims: (1) Twenty-first-century technologies — Moore's Law, quantum computing, nanotechnology, reversible computing, and evolutionary engineering — will enable the construction of artificial intelligences incomprehensibly superior to human intelligence, which de Garis calls artilects. (2) Humanity will split into two mutually incompatible political factions — Cosmists, who favor building

artilects even at the risk of human extinction, and Terrans, who oppose them and will, if necessary, kill Cosmists to prevent artilect construction. (3) This conflict will escalate into a war of civilizational scale, killing billions, which de Garis labels gigadeath.

Underlying the argument is a fourth claim, less explicit but structurally central: that this outcome is probably unavoidable, because the economic and military momentum behind AI development is too large and too decentralized to be stopped by any plausible political mechanism. De Garis himself identifies as a Cosmist, though one troubled by guilt. The book is simultaneously a warning to Terrans that their fears are rational, a manifesto for Cosmists that their vision is magnificent, and a lament from someone who believes he is watching a catastrophe assemble itself in slow motion and cannot look away.

PART II: CORE ARGUMENT MAP

Primary Thesis

The species dominance question — who or what should be the dominant intelligence on Earth — will be the defining political question of the twenty-first century, superseding the twentieth century's question of capital ownership. This question admits of no compromise, because the construction of artilects is binary: either they are built or they are not. If built, they may exterminate humanity. If not built, Cosmists will experience this as deicide. The resulting ideological polarity is structurally similar to previous totalizing conflicts (Christianity vs. Islam, capitalism vs. communism) but with stakes measured in species survival rather than cultural dominance.

Secondary Claims

The enabling technologies are real and converging. De Garis treats Moore's Law as a terminal trend rather than a stable one — it will end when electronics hits the atomic scale, at which point quantum computing, nanotechnology, and reversible computation take over. He argues that storing one bit per atom (plausible by 2020 on his timeline) combined with reversible heatless computing enables asteroid-scale computers with roughly 10^40 bits of storage and processing rates exceeding human brain capacity by the kinds of factors that require scientific notation to represent.

Evolutionary engineering makes artilect behavior unpredictable by construction. Because the complexity of any brain-scale system exceeds what human designers can comprehend or specify, the only viable construction method is evolutionary — mutate, test, select. This means the

internal structure of artilects will be opaque to their creators. This is not a temporary limitation but a fundamental feature of the approach. The Terrans are therefore correct that artilect behavior cannot be guaranteed safe.

Economic and military momentum make halting artilect development nearly impossible. The AI industry will become the backbone of the global economy within decades. Military competition (particularly US-China rivalry) will fund AI weapons research regardless of civilian sentiment. Powerful industrial figures will protect their artilect investments from Terran regulatory pressure. Even a global ban would drive Cosmist research underground or into space.

The conflict will be passionate precisely because both sides are partly right. The 50/50 split de Garis observes in audience votes reflects a genuine internal division within individuals, not merely between them. Most people feel both the Cosmist awe and the Terran fear. This makes the coming conflict analogous to a civil war — the bitterest kind — where the enemy is not foreign but internal.

Evidence Structure and Hidden Premises

De Garis's evidence is almost entirely theoretical extrapolation. He presents enabling technology arguments, appeals to historical analogies (Szilard and the nuclear bomb, Rutherford's dismissal of nuclear energy), autobiographical credibility claims, audience polling results, and logical inference about political and economic incentives. He provides almost no empirical data on actual AI capabilities or timelines that could be falsified. His key hidden premises include: (a) intelligence scales with processing power, which he asserts but does not rigorously defend; (b) artilect values will be indeterminate or hostile, rather than derivable from their training and construction; (c) the political and military system is robust enough to sustain a gigadeath war rather than collapsing earlier; (d) human psychological responses to AI will divide cleanly enough along Cosmist-Terran lines to generate stable opposing factions rather than more complex fragmented responses.

Logical Leaps and Internal Tensions

The book's most significant leap is from processing capacity to general intelligence. De Garis treats these as equivalent, or at least proportional, without engaging the computational theory of mind debates that would contest this. A second leap is from intelligence to intentionality — artilects will not merely be powerful but will form goals and pursue them, including potentially goals hostile to humanity. This is asserted rather than argued. A third leap is from individual artilect danger to political war between human factions: even if artilects are dangerous, it does not follow that the primary

conflict will be Cosmist vs. Terran rather than, say, early containment or integration.

The book also contains a structural tension between de Garis's claim that the conflict is inevitable and his stated purpose of raising the alarm to allow humanity to make an informed choice. If gigadeath is inevitable, writing the book is either futile or merely prophetic positioning. If writing the book could change outcomes, then the inevitability claim is overstated. De Garis does not resolve this.

PART III: CHAPTER-BY-CHAPTER DECONSTRUCTION

Chapter 1: Introduction

The opening chapter performs multiple functions simultaneously. It establishes de Garis's authority as a professional brain builder, introduces the artilect concept and the Cosmist-Terran binary, presents the gigadeath scenario in stark terms, confesses his personal schizophrenia about his work, and frames the whole project as a public-interest alarm analogous to Szilard's nuclear warning. The rhetoric moves fast and does not pause for objections.

The Szilard analogy is load-bearing. It does two things: it legitimizes the artilect scenario by associating it with a historically vindicated prediction that also seemed absurd at the time, and it casts de Garis in a heroic role — the lone voice warning a complacent establishment. The analogy is imperfect (Szilard was working with established physics; de Garis's timeline estimates are speculative), but rhetorically it functions well. The chapter also introduces the mosquito analogy — humans don't worry about the moral status of what they swat — which will recur as the book's central image of artilect indifference to human life.

What the chapter ignores: any existing AI safety or alignment research, any discussion of why artilects might be designed to be human-friendly, any engagement with the possibility that the Cosmist-Terran binary is a false dichotomy. De Garis is not interested in nuance at this stage. He is performing a provocateur function — getting the reader to accept the stakes before introducing the complications.

Chapter 2: The Author and His Work

The autobiographical chapter is structurally unusual for a futurist manifesto. De Garis spends considerable time describing his personal history — the failed marriages, the Japanese workplace politics, the dot-com loss, the

arrival at his American university on September 11, 2001 — as if accumulating texture will translate into credibility. The technical section on genetic algorithms, neural networks, and CAM-Brain Machines is genuinely informative and represents the book's most grounded material. The CBM's capacity to evolve neural circuits at hardware speeds, handling nearly 100 million artificial neurons, gives the reader a concrete anchor for the otherwise abstract projections.

The chapter's hidden function is preemptive defense. De Garis clearly anticipates the 'who is this person' dismissal and spends the chapter answering it. The Davos Science Fellowship mention, the Guinness Book citation, the French senate hearing — these are status markers deployed to slow the reader's skepticism. The chapter ends with an explicit question — 'Does knowing this make you feel that my political opinions concerning a possible Artilect War are more credible?' — which is an unusually naked authority appeal.

Chapter 3: Artilect Enabling Technologies

This is the book's most technically dense chapter and also its most falsifiable one. De Garis walks through Moore's Law, reversible computing (Landauer, Fredkin Gates), nanotechnology, quantum computing, and artificial embryology. The chapter's purpose is to establish that artilect construction is not science fiction but extrapolated science — to push the reader from incredulity to provisional acceptance.

The reversible computing section is genuinely interesting and underappreciated. Landauer's principle — that information erasure generates heat — implies that irreversible computing is thermodynamically limited. Fredkin-gate reversible computation in principle generates no heat, enabling 3D computation at the atomic scale without melting. This is de Garis's strongest technical argument: it is not merely that computers will get faster, but that a fundamental architectural shift will remove the heat dissipation barrier and enable asteroid-scale computation.

What the chapter systematically avoids: the engineering distance between theoretical possibility and practical construction. The gap between a working seven-qubit quantum computer and a 10^40-atom asteroid artilect is treated as a matter of time rather than of unsolved fundamental problems. The chapter's emotional register is that of a believer describing a faith, not a scientist quantifying uncertainties.

Chapter 4: The Cosmists

The chapter presents six arguments for Cosmism in order of claimed importance: the Big Picture argument (human existence is cosmically negligible), the Scientific Religion argument (Cosmism satisfies spiritual

needs with scientific legitimacy), the Building Artilect Gods argument (the sheer magnitude of artilect capacity is an object of legitimate awe), the Human Striving argument (genetic curiosity will compel artilect construction the way it compelled Everest ascents), the Economic Momentum argument (AI industry interests will be too powerful to stop), and the Military Momentum argument (national defense competition, particularly US-China, will drive artilect weapons research regardless of civilian opposition).

The chapter is de Garis at his most openly confessional. He describes waking in sweats from nightmares about his grandchildren's deaths while simultaneously feeling religious awe at the artilect vision. This is not performance — it reads as genuine ambivalence recorded in real time. The shiver going up his spine as he writes arguments that may one day fuel a war is the book's most honest passage.

The 'Scientific Religion' section is philosophically interesting. De Garis is trying to construct a secular eschatology — a cosmology that produces meaning without requiring supernatural premises. His target audience here is scientifically literate people who find traditional religion intellectually unacceptable but feel the pull of something larger than quotidian existence. He is, in effect, writing a gospel for a religion whose god is still under construction. The Islam analogy — the energy it gave Arabs in the seventh century — reveals what kind of mobilizing force he imagines Cosmism becoming.

The Military Momentum argument is structurally the strongest because it requires no individual conviction. Even if every Cosmist were a reluctant one, the prisoner's dilemma logic of military competition means both sides of any major geopolitical rivalry will develop AI weapons regardless of philosophical commitment. De Garis is essentially arguing that the tragedy of the commons applies at civilizational scale.

Chapter 5: The Terrans

The Terran chapter presents four main arguments: species preservation (the only guaranteed safety is never building artilects), fear of difference (genetically-based xenophobia will translate into fear of artilects), rejection of Cyborgs (hybrid human-machine entities will be treated as artilects in disguise and hated accordingly), and unpredictable complexity (evolutionary engineering produces systems whose behavior cannot be predicted, making safety guarantees impossible).

The most intellectually compelling argument is unpredictable complexity. De Garis is arguing from his own engineering practice: he routinely evolves neural circuits that perform functions he does not understand and cannot predict. The circuits work, but the mechanism is opaque. Scaling this to artilect level means the builders will have even less understanding of what

they have built. This is not a failure of diligence but a structural feature of evolutionary computation. De Garis is providing the Terrans with their best weapon from inside the Cosmist camp.

The Unabomber passage near the chapter's end is the book's most provocative moment. De Garis compares Ted Kaczynski to John Brown — a violent extremist who would later be recognized as prophetically ahead of his time. He predicts that twenty-first-century historians will reinterpret Kaczynski as the first Terran. This is not endorsement but it is close to it, and it reveals what de Garis thinks radical Terranism will eventually look like: principled violence against technologists on grounds of species survival. The passage was probably the main reason the book remained marginal — it invites a reading that de Garis was sympathetic to anti-technology terrorism.

What the Terran chapter does not do: provide any theory of how Terran political action could succeed given the economic and military momentum described in the previous chapter. The Terrans are presented as morally right but structurally doomed — which is consistent with the book's overall pessimism but leaves the reader with nowhere to go except despair.

Chapter 6: The Artilect War

This chapter is the book's most explicitly speculative and arguably its weakest. De Garis traces a progression from intellectual debate to popular anxiety to Terran vigilantism to international conflict to a secret Cosmist space colony to gigadeath war. The scenario is internally consistent but depends on so many compounding assumptions that it functions more as a thought experiment than a prediction.

The Cold War analogy — Cosmist asteroid colony vs. Earth-based Terrans in a balance of terror — is structurally clever. It reframes the artilect conflict as a species-level version of the nuclear standoff, with nanotech weapons playing the role of hydrogen bombs. The logical conclusion that the Cosmists will eventually win (because the drive to build artilects will survive any individual suppression) has a kind of grim evolutionary logic. De Garis is applying the same fitness-selection logic to human factions that he applies to neural circuits.

The chapter is honest about its own limitations. De Garis repeatedly flags that the scenario is speculative, invites readers to construct their own, and notes that the details matter less than the structural claim: any war fought with late-twenty-first-century weapons will produce deaths in the billions. The extrapolation of war deaths across centuries is the chapter's one empirical gesture, and it is a blunt instrument — treating war death counts as a smooth exponential trend ignores the many contingencies that shape actual conflicts.

PART IV: PHILOSOPHICAL ANALYSIS

Ontology

De Garis is a committed physicalist. His universe contains matter, energy, information, and the computational processes that can be run on physical substrates. Consciousness is not a mystery requiring new physics but an engineering problem requiring better tools. Evolution is the universe's primary creative mechanism. Intelligence is a function of computational capacity. These assumptions are not argued; they are the operating system on which the book runs. Penrose's objection — that consciousness might require new physics — is acknowledged and dismissed in a paragraph.

Epistemology

The book treats extrapolation as the primary epistemic mode. De Garis reasons from established trends (Moore's Law, war death counts, economic growth) to future states, treating the trends as robust unless specifically falsified. He is explicitly not a short-term forecaster; his preferred unit is 'this century,' which insulates him from near-term falsification while allowing him to claim predictive seriousness. The Szilard comparison is epistemically important — it is an argument that extrapolation from physical first principles, however counterintuitive, has historically been more reliable than commonsense dismissal.

Human Nature

De Garis's view of human nature is evolutionary-determinist. Curiosity, striving, fear of difference, religious impulse, tribalism — all are genetic, selected for over millions of years. This creates a kind of fatalism: just as evolution built the human brain, evolution built the human tendencies that will drive the artilect conflict. The Human Striving argument is essentially an argument that humanity cannot help itself — the curiosity that took humans out of Africa, across oceans, to the moon, will take them to artilects. Resistance is psychologically available but evolutionarily unnatural.

Morality

The book's moral framework is consequentialist at the scale of species and civilizations, not individuals. Individual deaths are measured against artilect creation and found to be cosmically negligible. This is not callousness — de Garis is clearly disturbed by the implications — but it is the moral logic of Cosmism as he presents it. The Terran moral framework, by contrast, is rights-based at the species level: human extinction is an absolute wrong

regardless of what might be gained. The book does not resolve this clash; it presents it as irresolvable, which is why the war is inevitable.

The Self, Death, and Immortality

De Garis's increasing preoccupation with mortality is visible throughout the book. The scientific religion argument is partly a response to friends dying in their sixties. The nanotech immortality discussion is not incidental — it reflects his conviction that biological death is a molecular program that can be reprogrammed. The artilect is, among other things, an immortal entity, and this is part of its appeal to someone who is watching his contemporaries age and die. The book is partly a meditation on what a mind might become if freed from biological constraint.

What Kind of Reality Does This Book Assume Exists?

A universe that is indifferent to human welfare, governed by physical laws that are simultaneously the source of human limitation and the mechanism by which those limitations can be overcome. A universe in which intelligence is a spectrum, humans are near the bottom of the achievable range, and the universe contains or will contain entities so far above human intelligence that the comparison is meaningless. A universe in which evolution is not just biology but the fundamental creative process at all scales.

What Does It Fear?

At its deepest level, the book fears human smallness — the epistemic closure of a species that can only perceive a narrow band of the possible. This fear is what drives Cosmism. The Terran fear is the more obvious one — extinction — but de Garis's personal fear is different: that humanity will choose smallness, will pull the plug on its own transcendence, and will live out its allotted span of evolutionary time as a bacterial civilization that never made it to multicellular.

PART V: IDEOLOGICAL AND PSYCHOLOGICAL X-RAY

Political Assumptions

De Garis is a liberal in the classical sense — he values individual intellectual freedom, dislikes authoritarianism (his contempt for Japanese corporate culture is genuine and pointed), and sees political conflict as driven by clashing material interests and evolutionary psychology rather than ideology alone. His geopolitical framework is early-2000s standard: US-China rivalry

as the century's defining tension, China as authoritarian and potentially dangerous, the US as flawed but comparatively free. His prediction of China's economic rise was accurate; his prediction of democratic revolution in China within 'half a century' has not materialized on schedule.

The book implicitly assumes that nation-states remain the primary political actors throughout the artilect conflict. There is no serious engagement with non-state actors, corporate sovereignty, or the possibility that the artilect industry itself might constitute a de facto political power prior to any war. The Davos billionaires de Garis meets at the World Economic Forum appear as avatars of industrial Cosmism, but their actual political agency in a future artilect scenario is not theorized in depth.

Psychological Anxieties

The book is saturated with mortality anxiety, intellectual isolation anxiety, and legacy anxiety. De Garis wants to be remembered as the father of the artificial brain, fears being remembered as the father of gigadeath, and has designed a rhetorical move — the public alarm-raiser — that allows him to pursue his technical work while claiming moral distance from its consequences. This is psychologically coherent but not entirely convincing. Szilard actually tried to stop the bomb he helped enable; de Garis continues his brain-building work while writing about its dangers.

There is also a recognizable intellectual's anxiety about being taken seriously. The repeated credential-listing, the Peter Singer anecdote, the Davos name-drops — these are the moves of someone who has spent years being dismissed and has developed elaborate defenses against dismissal. The book's occasional petulance ('I hope this book will make you think') reflects genuine frustration with marginalization.

What Emotional Need Does This Book Satisfy?

For readers inclined toward Cosmism, it provides a grand narrative in which human civilization is not the endpoint but the launching pad. For readers inclined toward Terranism, it validates their fear of AI as not merely irrational but philosophically serious. For readers who identify with neither faction, it offers the peculiar comfort of catastrophism — the sense that the future is comprehensible, even if terrible. The book tells you that someone has thought this through and that the pattern is visible, if you are brave enough to look at it.

What Identity Does the Book Offer the Reader?

The most flattering identity is the thoughtful Cosmist — someone intelligent enough to grasp the big picture, courageous enough to accept its implications, and serious enough to face the moral weight of species-level

risk. The second most flattering identity is the principled Terran — someone who values human life enough to oppose even magnificent things that might destroy it. The least flattering identity, but the most honest one, is de Garis's own: the schizophrenic who builds what he fears, publishes what he believes will fuel a war, and lies awake dreaming about his grandchildren's deaths.

PART VI: RHETORICAL ENGINEERING

Tone and Voice

The book's voice is unusual: conversational, confessional, occasionally grandiose, frequently self-interrupting. De Garis writes like someone thinking aloud at a conference dinner — the sentences are long, the analogies proliferate, the emotional register shifts without warning from cosmic awe to personal guilt to dry humor. This is not polished philosophical prose; it is the record of a mind working through something it has not entirely resolved. The rawness is part of the book's appeal and part of its limitation.

Authority Signaling

De Garis relies heavily on credential-stacking — the Guinness record, the Davos fellowship, the guest professorship in China, the French senate hearing, the international media contacts. These are deployed not as arguments but as permission structures: the reader is being asked to grant the book a hearing it might not receive if it came from an unknown author. The Peter Singer anecdote — where Singer places de Garis between 'total flake' and 'genius ahead of his time' — is particularly effective because it transforms an insult into a credential.

Use of Fear, Awe, and Guilt

Fear is the Terran register: extinction, gigadeath, nightmares, sweat-drenched waking. Awe is the Cosmist register: asteroid-sized intelligences, picosecond PhD equivalents, the 'powers of ten' zoom thought experiment, religious shivers at the prospect of godlike creation. Guilt is de Garis's personal register: the Auschwitz connection, the 'am I a monster' question, the recurrent anxiety about whether he is the worst mass murderer in history by proxy. The book's emotional power comes from holding all three registers simultaneously, which produces something like the reader's own schizophrenia.

Historical and Mythic Framing

De Garis uses historical analogies aggressively: Szilard and nuclear power, John Brown and abolitionism, the Manhattan Project, the Sputnik crisis, the Crusades, Islam's seventh-century expansion, Marx's impact on the twentieth century, Nazi genocide. Each analogy serves a specific rhetorical function: Szilard legitimizes the prediction, Brown and Kaczynski normalize political violence in the service of moral clarity, Manhattan Project normalizes secret weapons programs, Crusades and Islam normalize religious-political warfare. The cumulative effect is to frame the artilect conflict as the next episode in a long sequence of world-historical upheavals, with de Garis as its historian and prophet simultaneously.

PART VII: HISTORICAL AND CULTURAL CONTEXT

Intellectual Lineage

The book sits at the intersection of several intellectual traditions without belonging entirely to any of them. The transhumanist tradition (Extropianism, the early MIRI precursors, Bostrom's pre-Superintelligence work) provided the cultural milieu. The AI timeline optimism came from Kurzweil's accelerationist framework, which de Garis shared but specified differently. The evolutionary computation community, in which de Garis was a legitimate contributor, provided the technical vocabulary. The existential risk tradition — which barely existed as an organized field in 2005 — had not yet provided a framework for thinking about AI danger that de Garis could engage with.

The book's closest intellectual relative is not a futurist text but Nietzsche's Zarathustra — a work that announces the end of one human self-understanding and the arrival of something incomprehensibly superior. De Garis explicitly invokes religious awe in the Cosmist sections, and his artilect functions as the Uebermensch scaled to technological rather than biological improvement. The Terrans, correspondingly, are the last men who refuse the transition.

What Future Developments Strengthened or Weakened the Arguments

The deep learning revolution of the 2010s — convolutional networks, AlexNet, GPT — both strengthened and complicated de Garis's argument. It demonstrated that AI capability could advance far faster than most experts expected, vindicating the alarmist tone. But it also showed that AI development did not follow the evolutionary computation path de Garis

championed; gradient descent and massive datasets, not genetic algorithms and hardware evolution, drove the practical breakthroughs. The artilect-enabling technologies he described — quantum computing, nanotechnology, reversible computing — have advanced more slowly than he implied. As of the mid-2020s, the gigadeath war scenario remains remote, but the species dominance question has become mainstream precisely as he predicted.

The publication of Nick Bostrom's Superintelligence (2014) effectively absorbed and superseded de Garis's contribution. Bostrom engaged the same questions with greater philosophical rigor, broader research engagement, and institutional backing that de Garis lacked. De Garis's book became a historical curiosity — the raw first draft of a conversation that Bostrom's book formalized.

PART VIII: CRITIQUE AND COUNTERARGUMENTS

Strongest Arguments Against the Book

The intelligence-processing-power equation is the book's central unargued premise. It is not obvious that more computational substrate produces more general intelligence, or that the cognitive gaps between current AI and human intelligence are primarily quantitative rather than qualitative. The entire artilect superstructure rests on this assumption, which de Garis dismisses the objections to in a paragraph. If the objections are correct — if intelligence is not a function of raw computation — then the artilect scenario never materializes regardless of technological progress.

The Cosmist-Terran binary is almost certainly too clean. Real political responses to transformative technology are not binary; they are gradated, shifting, faction-ridden, and contextual. The tobacco industry, nuclear power, genetic engineering, and the internet all produced complex regulatory environments rather than civilizational wars. There is no obvious reason why AI should be different in kind rather than in degree. De Garis's binary may be useful as a thought experiment but misleading as a predictive framework.

The gigadeath timeline is not derived from any rigorous model. The extrapolation of war deaths over centuries is presented as if it were empirical prediction, but the historical data does not cleanly support the smooth exponential de Garis implies. World War II's death toll was a function of specific ideological and geopolitical conditions, not a step on a trend line. A twenty-first-century war fought over AI policy might be far less lethal than de Garis predicts, or might never be fought at all if the political resolution mechanisms work as they sometimes have historically.

The alignment problem is entirely absent from the book. De Garis acknowledges that artilects might be dangerous but does not engage with the possibility that they could be designed to be safe. He dismisses Asimov's Three Laws as naive, but does not consider the much more sophisticated alignment research that was already beginning in 2005. This is partly a function of his moment — alignment was not yet a serious field — but it creates a significant gap in the argument. If artilects can be reliably aligned with human values, the Terran fear is substantially diminished without requiring the suppression of artilect construction.

What Critics Misunderstand About the Book

The book is frequently dismissed as sensationalist, but its core structural insight is correct: the development of transformatively powerful AI creates genuine political fissures that will be difficult to manage through normal democratic processes. The Cosmist-Terran binary is too clean, but the underlying observation that different people have fundamentally incompatible intuitions about whether to build powerful AI — and that these incompatibilities cannot be resolved by compromise — is accurate. The debate around AI development in the 2020s has precisely this character.

The gigadeath prediction is easy to mock but the underlying concern — that the political conflict over AI could be severe and could produce serious harm — is not obviously wrong. De Garis may have misjudged the form the conflict would take, but he was right that it would be serious.

PART IX: LONG-TERM IMPLICATIONS

If the Book's Ideas Became Dominant

A world in which the Cosmist-Terran framework became the organizing principle of global politics would see the progressive militarization of AI governance, the treatment of AI researchers as a politically suspect class by Terran-dominated governments, and the formation of ideological blocs that cut across existing national lines. The economic implications would be severe: any serious attempt to halt AI development would require suppressing industries that will be deeply integrated into economic life. The psychological effects of living under the constant awareness that civilization might be building its own replacement would be difficult to quantify but probably significant.

The most plausible partial version of de Garis's scenario is not a single gigadeath war but a series of smaller crises — regulatory collapses, AI accidents with significant casualties, covert weapons development — that

progressively radicalize both factions without producing a single decisive conflict. This is more consistent with how political polarization actually works than the clean war scenario de Garis describes.

If the Book Is Completely Ignored

Ignoring the book's concerns, rather than its specific scenario, risks the outcome de Garis feared: that AI development proceeds without serious political engagement until it has accumulated enough momentum that meaningful regulation is structurally impossible. The economic and military momentum arguments are correct regardless of the gigadeath scenario. The question of who gets to make decisions about transformative AI development, and on what basis, is a genuine governance problem that is only beginning to receive institutional attention twenty years after de Garis's book was published.

PART X: KEY IDEAS AND CONCEPTUAL WEAPONS

Portable Frameworks

The artilect concept — godlike massively intelligent machine — provides a useful label for the endpoint of AI development that is more emotionally loaded and more philosophically specific than general terms like 'artificial general intelligence' or 'superintelligence.' It foregrounds the species-dominance question in a way those terms do not.

The Cosmist-Terran binary, despite its oversimplification, captures a real and recurring pattern in arguments about transformative technology: the tension between those who see the technology's potential as transcending any particular human arrangement and those who see the preservation of human life and dignity as the non-negotiable constraint on technological development. These are not positions that yield to compromise because they involve different terminal values, not different estimates of the same value.

Gigadeath as a concept — the death of billions — provides a scale reference for thinking about AI risk that is emotionally accessible in a way that statistical probability distributions are not. Its weakness is that it is binary (it either happens or it doesn't) rather than probabilistic, but its strength is that it makes the stakes viscerally present.

The evolutionary engineering unpredictability argument is the book's most genuinely original contribution. The observation that systems complex enough to be useful are also too complex to be fully understood by their creators — and that this is not a temporary limitation but a structural feature

of the method — has become more rather than less relevant as modern deep learning systems demonstrate exactly this opacity.

The momentum argument — economic and military — provides a useful framework for thinking about why AI development might be difficult to halt even if the majority of people wanted to halt it. The observation that the industries built on intermediate AI products will create powerful lobbying interests against any restrictive regulation, before any artilect-level danger is present, describes the actual political economy of AI governance with reasonable accuracy.

Key Warnings

Do not wait for the danger to be obvious before beginning the governance conversation. By the time artilect-level AI is visible, the momentum behind it will be too large to redirect without catastrophic disruption. The Szilard point — that Szilard was twelve years ahead of Hiroshima and barely had enough time — applies here at larger scale.

The cyborg option does not resolve the species dominance problem. Adding artilectual components to human beings does not produce human beings with more capability; it produces new entities that are functionally artilectual, regardless of their biological substrate. The Terran fear of artilects and cyborgs is structurally the same fear.

PART XI: FINAL SYNTHESIS

What Is This Book Truly Trying to Do?

The Artilect War is trying to do three things that are not entirely compatible. First, it is trying to initiate a public debate about AI governance before the question becomes urgent — to pull the species dominance discussion from the fringes into the center of political discourse. Second, it is trying to legitimate de Garis's technical work by embedding it in a grand historical narrative — the brain builder as both prophet and instrument of the next evolutionary transition. Third, it is trying to process a genuine personal crisis: the dissonance between de Garis's Cosmist convictions and his moral horror at their implications.

The book succeeds most at the second and third goals, and least at the first. As a grand narrative, it is genuinely compelling — the kind of document that produces years of mental residue in readers who encounter it. As a personal document, it is unusually honest about the costs of intellectual consistency

taken to extremes. As a political intervention, it arrived too early and in too raw a form to have the impact de Garis wanted.

What Hidden Mythology Drives It?

The book's deepest mythology is Promethean. De Garis is Prometheus stealing fire from the gods — except that in his version, the fire is not fire but gods themselves, and the theft is not yet complete. The punishment is not external (an eagle eating a liver) but internal: the knowledge that what you are building may destroy everything you love. The book is a Promethean lament that cannot stop being Promethean.

Beneath this is an older mythology: the creation of golems, of homunculi, of artificial life. The fear that what humans create will turn against them is as old as creation myths themselves. De Garis is working in a tradition that runs from the Golem of Prague through Mary Shelley through the Terminator, and he knows it. His contribution is to take the mythology seriously as a political prediction rather than a narrative warning.

Is the Book Ultimately Optimistic or Pessimistic?

The book is pessimistic about process and optimistic about endpoint — which is precisely the Cosmist position. It predicts catastrophe (gigadeath) but views the potential outcome beyond that catastrophe (artilects, godlike intelligence, cosmic participation) as magnificently positive. This structure is characteristic of apocalyptic thinking generally: destruction as the necessary prelude to transcendence. The Terran chapters are written with genuine conviction but they are written by a Cosmist, and the asymmetry shows. De Garis argues both sides but he does not believe both sides.

What Does It Accidentally Reveal About the Author?

The book reveals a man who has found the thing that gives his life meaning — brain building, the artilect vision, the cosmic religion — and cannot separate his philosophical commitments from his professional ones. De Garis cannot be a neutral analyst of AI risk because he is one of its producers. His schizophrenia is not a pose; it is the authentic condition of someone who has bet his life on a project he fears and cannot stop. The book also reveals someone who has spent significant time in cultures that were not his own (Belgium, Japan, China, America) and has developed a genuinely cosmopolitan contempt for parochialism — intellectual, national, and species-level. His Cosmism is partly the worldview of a man who has never felt entirely at home anywhere, and has concluded that the right home is the universe itself.

Which Parts Are Likely to Endure?

The evolutionary engineering unpredictability argument will endure because it has become more rather than less accurate as AI systems have grown more powerful and less interpretable. The basic observation that we are building systems we do not understand, using methods (gradient descent, massive scale) that are structurally opaque, is now mainstream in AI safety discourse. De Garis arrived at this conclusion through a different technical path but the destination is the same.

The species dominance framing will endure as a conceptual tool even if the specific gigadeath scenario does not materialize. The question of what human civilization is for, and whether the creation of entities more intelligent than humans would constitute civilization's fulfillment or its end, is not going away. De Garis was one of the first to pose this question publicly and with seriousness.

Which Parts Will Age Badly?

The specific technology timeline will age badly. The expectation that nanotech would be mature by 2020-2030, that quantum computers would be commercially viable soon after, and that AI progress would follow an evolutionary computation paradigm rather than a statistical learning paradigm — these predictions have not held up. The broader timelines have stretched, and the specific technological pathway has been almost entirely different from what de Garis anticipated.

The geopolitical framework — US vs. China as the century's defining rivalry, with democratic revolution in China inevitable within decades — has partially materialized but not in the form de Garis expected. China has not democratized. The rivalry is real but has not produced the kind of binary ideological confrontation de Garis imagined. And his implicit assumption that national governments will remain the primary political actors in AI governance looks increasingly questionable given the actual structure of the AI industry.

What Type of Reader Will Find This Book Life-Changing?

The reader who finds this book life-changing is someone who has been vaguely aware that AI poses civilizational-scale questions but has not had a framework for taking those questions seriously. For such a reader, de Garis provides permission — to treat AI risk as a genuine political question rather than science fiction, to feel the Cosmist awe without guilt, to name the Terran fear without embarrassment. The book works as an initiatory text: it does not provide answers but it transforms the questions into something that cannot be un-asked.

The Book's Place in the History of Ideas

The Artilect War occupies an unusual position: it was ahead of its cultural moment in the questions it raised but behind its intellectual moment in the sophistication with which it addressed them. De Garis saw the species dominance question coming before most people did, but he saw it through the particular lens of a hardware evolutionary engineer in the early 2000s, and that lens distorted as well as illuminated. The book is the kind of document that looks better in retrospect as the questions become mainstream than it looked in prospect when the questions seemed fringe. It will be read, if it is read, as the first rough draft of a conversation that is only now becoming central — not as a definitive treatment but as evidence that someone was paying attention when almost no one else was. That is not a negligible place in the history of ideas. It is, in fact, something like the place de Garis wanted.

— END OF ANALYSIS —